arXiv:1409.6911v3 [cs.CV] 18 Nov 2014

Do More Drops in Pool$_5$ Feature Maps for Better Object Detection

Zhiqiang Shen
Fudan University
zhiqiangshen13@fudan.edu.cn

Xiangyang Xue
Fudan University
xyxue@fudan.edu.cn

Abstract

Deep Convolutional Neural Networks (CNNs) have gained great success in image classification and object detection. In these fields, the outputs of all layers of CNNs are usually considered as a high dimensional feature vector extracted from an input image and the correspondence between finer level feature vectors and concepts that the input image contains is all-important. However, fewer studies focus on this deserving issue. On considering the correspondence, we propose a novel approach which generates an edited version for each original CNN feature vector by applying the maximum entropy principle to abandon particular vectors. These selected vectors correspond to the unfriendly concepts in each image category. The classifier trained from merged feature sets can significantly improve model generalization of individual categories when training data is limited. The experimental results for classification-based object detection on canonical datasets including VOC 2007 (60.1%), 2010 (56.4%) and 2012 (56.3%) show obvious improvement in mean average precision (mAP) with simple linear support vector machines.

1. Introduction

Object detection is a fundamental and crucial problem in computer vision. One of the most heavily studied paradigms and the most prominent example for object detection is deformable part-based models (DPM) algorithm [9]. It combines a set of discriminatively trained parts in a star model which is called pictorial structure [13, 10, 11]. The part filters in DPM are based on hand-crafted Histogram of Gradients descriptors [3]. However, the progress has been slow during 2010-2012 in the canonical visual recognition task PASCAL VOC object detection [6] with hand-crafted visual features.

In the last years, more and more works focus on Deep Convolutional Neural Networks (CNNs) and achieve great success. CNNs were firstly introduced in 1980 by Kunihiko Fukushima [14], and Yann LeCun et al. improved it in 1998 [25]. This model was initially applied to handwritten digit recognition [24] and OCR [25]. Recently CNNs have been well applied into lots of visual recognition systems in a variety of domains. With the introduction of large labeled image databases [4] and the massive parallel computations of GPU implementations, the large scale CNNs have become the most accurate method for generic object classification [23] , detection [15, 29, 31] and segmentation tasks [17, 16].

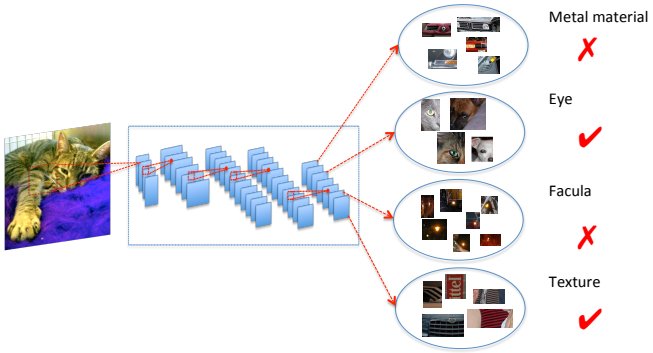

Figure 1. Illustration showing the main motivation of this paper. For a parameter-trained network, correspondence between pool$_5$ features and concepts of each channel would also be fixed, such as shape, texture and material properties. It is clear that some concepts are helpful and some are unfriendly for classification.

In 2012, Krizhevsky et al. [23] designed a large CNN with 60 million parameters and 650,000 neurons and obtained substantial improvement of image classification accuracy on the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [22, 4]. In 2013, a joint deep learning architecture was proposed for pedestrian detection, which combines a CNN architecture with the DPM algorithm [27]. Four components are contained in this framework: feature extraction, part deformation handling, occlusion handling, and classification. In 2014, Girshick et al. [15] proposed a scalable detection algorithm called R-CNN and showed that R-CNN can obtain dramatically higher object detection performance on PASCAL VOC as compared to algorithms based on HOG-like features. R-CNN is a region-based al-

gorithm that bridges the gap between image classification and object detection by operating within the "detection with regions" paradigm [32] instead of the sliding-window strategy.

For further improving detection performance, several methods were studied before. One kind approach is to manipulate the training images by different operations, such as rotation, scaling, shearing and stretching, and then merge these transformed images into training set for training a more powerful detector which will improve the view-invariant representations [20]. Another kind of approach is to perform local transformations in feature learning algorithms like Restricted Boltzmann Machine (RBM) and autoencoders, which combines various kinds of transformed weights (filters) and expands the feature representations with transformation-invariance [30]. In reality, occlusion and inconsistent property such as the same crowds or objects with different colors, texture or material properties often exists between training and testing data sets. So simply considering invariance is still far from enough.

Considering this deficiency, we hold the opinion that establishing correspondence on middle-level features with input images is required. Some works have paid attention to this idea [15, 26]. Ross Girshick et al. [15] visualize activation units at $5th$ convolutional layer in the middle of each channel after pooling operation (the pool_5 feature map is $6\times6\times256$ dimensional) and find that units in pool_5 somewhat characterize concepts of objects (people or text), or texture and material properties, such as dot arrays and specular reflections. If we dropout some activation units to zeros, it seems that we perform some changes in the input images of CNN. If activation units with large intra-class and small inter-class variations set to zeros in the training phase, what would happen for object detection? Fortunately our experimental results give us positive answers.

Inspired by this, we propose a feature edit algorithm by finding out the distribution of concepts that pool_5 feature units correspond to. Our method is an entropy-based model which is to compute the probability that each concept owns in training data and drop out some lowest ones. That is to say we drop out those units unfriendly to classification in the pool_5 feature maps of CNN. Our algorithm named as $FeatureEdit$ is different from feature selection algorithms mainly for dimensionality reduction in order to lower the computational complexity and improve generalization [5].

Automatic feature selection can be formulated as the problem of finding the best subset S of features from an initial, and maybe a very large set of features $F(i.e., S \subset F)$. Since the ultimate goal is to improve performance, one could define the optimal subset of features as which provides the best classification ability in the given task. We measure it by a criterion function G as $c = G(S, D, M)$ where value c denotes the model classification ability, D denotes the data set used and M denotes the model parameters applied in the task. Our proposed feature edit algorithm belongs to the so-called $filter$ methods which could be heuristically defined and the superiority would be evaluated by the learning algorithms independent of the classifiers. The criterion function of our algorithm can be formulated as $c = G(E,\ D,\ M)$ where E is the edited version of F (i.e. feature sets are extracted from all training images by CNN). We use $5th$ convolutional layer features after max pooling operations for editing but still map to $7th$ layer's feature maps by multiplying the weights of w_6 and w_7 and then feed the new feature set to the boosted linear-SVM for training.

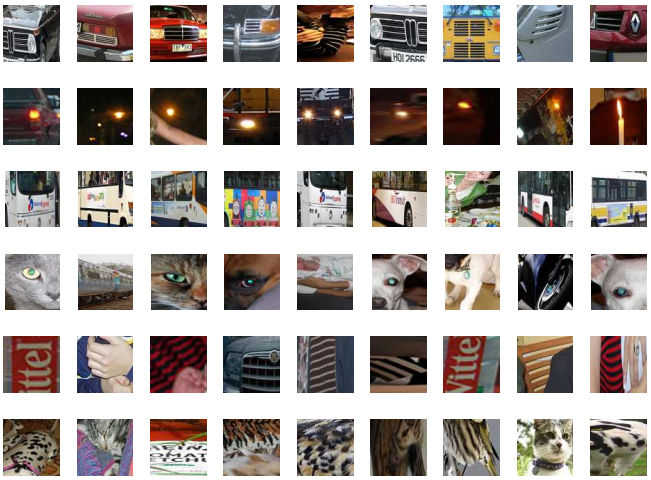

Figure 2. Illustration showing top 9 regions (receptive field) for six pool_5 units by a trained network. Some units are aligned to shapes, such as bus(row 3) or eye (row 4). Other units capture texture and material properties, such as dot arrays (row 6) and specular reflections (row 2).

This paper makes two important contributions: (1) We find out that we can easily obtain slightly modified CNN features by performing more dropouts in pool_5 without decreasing the ultimate object detection performance. (2) We present an entropy-based model to mark the subsets of pool_5 feature maps, which do not benefit for classification. We achieve obviously higher performance of 60.1%, 56.4% and 56.3% mAP on PASCAL VOC 2007, VOC 2010 and VOC 2012, which are higher than R-CNN and other competition approaches.

2. Concepts Correspondence

Establishing correspondence on a finer level than object category is required to understand convolutional neural networks. Jonathan Long et al. [26] studied the effectiveness of activation features for tasks requiring correspondence and present evidence that CNN features localize at

a much finer scale than their receptive field sizes. On the other hand, Zeiler and Fergus [35] introduced a novel visualization technique that gives insight into the function of intermediate feature layers, which could explain the impressive classification performance of the large convolutional neural networks. Each layer in classification CNN models shows the hierarchical nature of the features in the network. Here we focus on the detection networks which are slightly different from classification ones. Detection networks are fine-tuned from neat objects, but classification networks are trained on global images with large background area. In this paper, we just concentrate on the $5th$ layer features and how CNN models identify variant of objects.

Our concepts correspondence method follows the feature visualization in R-CNN [15] which sorts massive region proposals from highest to lowest activation. We make several slight modification in our experiments, one is the activation units index, the pool_5 feature map is $6 \times 6 \times 256$ dimensional (xIndex$\times$yIndex$\times$channel), we move the activation indexes from $(3, 3, channel)$ to $Max(3:4, 3:4, channel)$. Another modification is the region proposals, proposals in the original method are obtained by selective search. However, our proposals are randomly sheared from every object area. In each object, we shear about one thousand patches to order and perform non-maximum suppression for filtering.

Each row in Figure 2 shows the top 9 activations for a pool_5 unit from the CNN that we fine-tuned on VOC 2012 trainval. It implies that these units are aligned to shapes, textures, properties and other concepts. So if we set some particular channel units to zeros, the input image will be manipulated indirectly on concept-level.

3. Maximum Entropy Model

3.1. Preliminaries

Let $\mathcal{X} \subseteq \mathbb{R}^d$ and $\mathcal{Y}$ be the input and output space[1], $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$. We define a train sample as $(x_i, y_i), i = 1, 2, \ldots, 256$.

The correspondence between each channel feature vectors and concept space is defined by $kurtosis$ which measures peakedness of a distribution. $kurtosis$ is defined as

$$x(a) = \frac{E[o^4]}{E^2[o^2]} - 3$$

where $o^\alpha = (a_i - \bar{a})^\alpha$ and $E[\cdot]$ is the expectation operator for vector segments. a_i denotes the units in ith channel of pool_5 feature map. The -3 operator is to let normal-distribution kurtosis approach zero.

Let $\{f_t(x, y), t = 1, 2, \ldots, T\}$ be the feature function for different classes, which is a real-valued function. T denotes the category number in training dataset. $\tilde{p}(x, y)$ denotes the empirical distribution on training data. Feature function is defined as

$$f_t(x, y) = \begin{cases} 1 & if\ \tilde{p}(x, y) > threshold \\ 0 & else \end{cases}$$

The expectation of empirical distribution $\tilde{p}(x, y)$ is defined as

$$E_{\tilde{p}}(f_t) = \sum_{x,y} \tilde{p}(x, y) f_t(x, y)$$

The expectation of distribution $p(x, y)$ is defined as

$$E_p(f) = \sum_{x,y} \tilde{p}(x) p(y|x) f_t(x, y)$$

3.2. Problem Definition

Maximum entropy [1] can be used to estimate any probability distribution in its general formulation. In this paper we are interested in improving the performance to classify different objects. Thus we limit our further discussion to learn conditional distributions from labeled training data. The motivation of our idea is to find out the probability distribution that concepts are representative of the input object images and dropout those unfriendly ones for final classification.

Suppose that we are given T feature functions f_t for each class in training data. We would like our model to accord with the statistics that we train. Our goal is to extract a set of facts about the decision-making process from the training data that will aid us in constructing a model of this process. So we would like p to lie in the subset $\mathcal{C}$ of $\mathcal{P}$ which defined by

$$\mathcal{C} \equiv \left\{ p \in \mathcal{P} | E(f_t) = \tilde{E}(f_t) \quad for\ t \in \{1, 2, \cdots, T\} \right\}$$

The conditional entropy is defined as

$$H(p) = -\sum_{x,y} \tilde{p}(x) p(y|x) \ln p(y|x)$$

In maximum entropy, we use the training data to set constraints on the conditional distribution. To select a model from a set $\mathcal{C}$ of allowed probability distribution, choose the model $p^* \in \mathcal{C}$ with maximum entropy $H(p)$

$$p^* = \arg\max_{p \in \mathcal{C}} H(p)$$

4. Nonparametric Algorithm for p^*

Instead of referring to original constrained optimization problem as the $primal$ problem, we propose a nonparametric method to find out p^*. We call our method as $Edit$ algorithm which is based on the intra-class and inter-class variations of activation value x_i in each channel.

[1] Input space means the pool_5 feature map; Output space means the measured concept presentation for each channel.

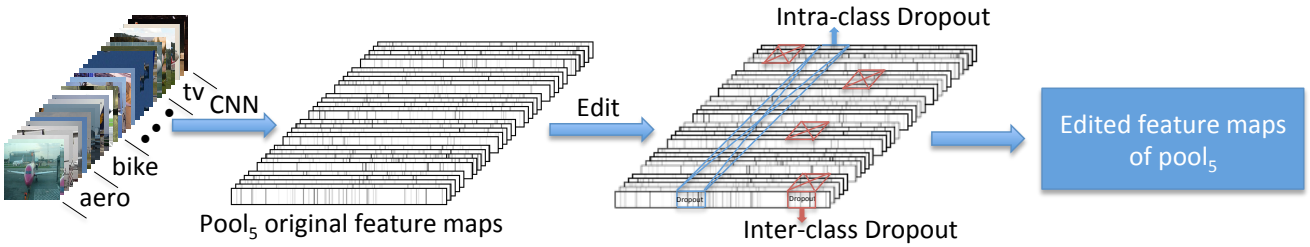

Figure 3. Illustration showing the edit procedure. From left to right: (1) Input regions (ground truth), which are resized to 227×227, (2) Pool_5 feature map extracted from CNN, pool_5 means the $5th$ convolutional layer features after max pooling operation, (3) Edit operation, (4) Outputs of the edited features.

4.1. Algorithm Details

- Intra-class Edit

  In this stage we would like to find out the subset of each channel features which has the largest intra-class variations.

  We define the training set $\{x_i\}$ as $K_{ji}^C(\cdot)$ which denotes the value of Cth class in jth training sample and ith channel segment, $i = 1, 2, \cdots, 256$, $j = 1, 2, \cdots, N_C$, $C = 1, 2, \cdots, T$, N_C is the number of training examples in each class C.

  Then compute the variance of each feature segment's statistic in class C which is defined as

$$V_i^C = \frac{1}{N_C} \sum_{j=1}^{N_C} (K_{ji}^C(a^i) - \bar{K}_{ji}^C(a^i))^2$$

  where V_i^C denotes the variance of Cth class training feature vectors in ith channel feature segment. The larger of V_i^C means channel i is unsteady and unfriendly for classification.

  Following it we compute intra-class p^* by V_i^C

$$p_i{}^*(intra) = \frac{V_i{}^C}{\sum_i V_i{}^C}$$

- Inter-class Edit

  We find out subsets with the smallest variations for all classes. First compute the mean value of the statistics at channel i in Cth class:

$$\bar{K}_i^C(a^i) = \frac{1}{N_C} \sum_{j=1}^{N_C} K_{ji}^C(a^i)$$

  Then compute the mean value of the statistics in all T classes:

$$\bar{K}_i^A(a^i) = \frac{1}{T} \sum_{C=1}^{T} \bar{K}_{ji}^C(a^i)$$

  where $\bar{K}_i^A(\cdot)$ denotes the average statistic of all classes.

  The variance of the statistics in all classes is defined as

$$V_i^A = \frac{1}{T} \sum_{C=1}^{T} (\bar{K}_i^C(a^i) - \bar{K}_i^A(a^i))^2$$

  Then we compute inter-class p^*

$$p_i{}^*(inter) = \frac{V_i{}^A}{\sum_i V_i{}^A}$$

For the original training feature set $F_{CNN} = \{F_1, F_2, \cdots, F_{N_C}\} \in \Re^{N_C \times k}$, where k is the length of one CNN feature vector. $F_i(i = 1, 2, \cdots, N_C)$ denotes one feature vector extracted from an image. The edited feature values are defined as

$$x_{edit} = x \otimes f(x, y)$$

where $f(x, y)$ is defined by p^* with 20% intra-class and 30% inter-class $threshold$. And then we can obtain F_{Edit} by x_{edit}. For each channel units in F_{CNN}, if $x_i = 0$, dropout all units in this channel to zeros. Pipeline is shown in Figure 3.

Up to now, we have obtained edited features F_{Edit} by dropping out selected channels in $pool_5$ feature maps according to the proposed algorithm. Because of the same dimensionality with original pool_5 feature maps F_{Ori}, F_{Edit} can feed directly to the fully-connected layers of fc_6 and fc_7.

5. Experiments and Discussions

We evaluate performance on the datasets: Pascal Visual Object Challenge (VOC) 2007 [7] and VOC 2010-2012 [8]. VOC 2007 dataset contains 5011 trainval (training + validation) images and 4952 test images over 20 classes. 10103 trainval images and 11635 test images are in VOC 2010. The data sets are identical for VOC 2011 and 2012 which both contain 11540 images in trainval set and 10991 images in test set.

VOC 2007	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
FE_R	69.1	72.7	54.9	42.0	35.0	66.1	74.1	64.9	36.8	67.7	55.1	63.6	67.3	**71.9**	57.7	32.0	65.5	51.6	64.1	64.9	58.8
FE_S	68.0	69.1	52.9	41.5	35.1	66.1	72.9	67.4	36.2	65.4	51.8	61.0	65.8	70.9	55.7	31.7	64.7	51.9	61.5	64.2	57.7
FE_E	**71.3**	71.6	56.1	42.2	37.1	67.2	74.4	67.9	36.6	68.2	54.3	64.7	70.2	70.9	58.9	34.7	**66.3**	**53.8**	64.2	67.3	59.8
FE_M	71.0	72.1	55.2	41.5	36.3	**69.5**	**74.7**	67.2	**37.2**	**68.6**	57.1	**64.7**	69.8	71.8	**59.1**	**35.3**	65.8	52.9	64.4	67.8	**60.1**
R-CNN	68.1	**72.8**	56.8	**43.0**	36.8	66.3	74.2	67.6	34.4	63.5	54.5	61.2	69.1	68.6	58.7	33.4	62.9	51.1	62.5	64.8	58.5
SPP	68.6	69.7	**57.1**	41.2	**40.5**	66.3	71.3	**72.5**	34.4	67.3	**61.7**	63.1	**71.0**	69.8	57.6	29.7	59.0	50.2	**65.2**	**68.0**	59.2
DNP+R											—										46.1
Regionlets	54.2	52.0	20.3	24.0	20.1	55.5	68.7	42.6	19.2	44.2	49.1	26.6	57.0	54.5	43.4	16.4	36.6	37.7	59.4	52.3	41.7
Szegedy	29.2	35.2	19.4	16.7	3.7	53.2	50.2	27.2	10.2	34.8	30.2	28.2	46.6	41.7	26.2	10.3	32.8	26.8	39.8	47.0	30.5
DPM v5	33.2	60.3	10.2	16.1	27.3	54.3	58.2	23.0	20.0	24.1	26.7	12.7	58.1	48.2	43.2	12.0	21.1	36.1	46.0	43.5	33.7
$\triangle_{R-CNN}$	+2.9	-0.7	-0.4	-1.5	-0.5	+3.2	+0.5	-0.4	+2.8	+5.1	+2.6	+3.5	+0.7	+3.2	+0.4	+1.9	+2.9	+1.8	+1.9	+3.0	+1.6

Table 1. Detection average precision(%) on PASCAL VOC 2007 test. Rows 1-4 show our experimental results. Feat_R: Random feature edit algorithm; Feat_S: Images shearing algorithm; Feat_E: Only using edited features for training; Feat_M: Merging the original features and edited features for training. Rows 5-10 show other competition approaches. (R-CNN [15]; SPP (without combination) [18]; DNP+R [37]; Regionlets [33]; Szegedy et al. [31]; DPM v5 [9]). Row 11 shows the differences between FE_M and R-CNN.

5.1. Training Stage

There are two parts in training stage: train parameters of convolutional neural network and train a linear classifier. Details are as follows.

5.1.1 Network Parameters Training

We use the Caffe [21] implementation of the CNN defined by Krizhevsky et al. [23], which is used in various kinds of domains such as fine-grained category detection [36] and object detection [15]. It consists of total seven layers, the first five are convolutional and the last two are fully connected. Our training strategy is supervised pre-training on a large auxiliary dataset (imagenet 2012 train-val dataset) and domain-specific fine-tuning on VOC 2012 trainval dataset. This strategy is popular in recent years such as R-CNN [15]. Although VOC 2012 is the largest in PASCAL VOC datasets, it is still not enough to fine-tune the parameters to a good location. Thanks to the improvement of generic objectness measures which produce a small set of candidate object windows from raw images, we select all candidate object proposals with $\geq$ 0.6 IoU overlap with a ground truth box to rich the fine-tuning dataset. Two existing high performance objectness detectors have been tried: BING [2] and selective search [32]. We find that BING is faster than selective search but the IoU overlap with ground truth is lower than the latter. In this paper we use selective search for pre-detecting, but if you care efficiency more, BING will be a better choice. CNN fine-tuning is run for 70k SGD iteration on VOC 2012 trainval dataset and selected windows. The CNN we used requires a fixed-size input of 227×227 pixels. R-CNN [15] has evaluated two approaches for transforming object proposals into CNN inputs and finds that warping with context padding (p=16 pixels) outperformed other approaches (more details in R-CNN). Finally a fixed-dimensional feature vector from each proposal is extracted. Our CNN parameters and algorithm codes will be released soon.

5.1.2 Linear-SVM and Regression

A linear classifier is trained on different feature sets respectively, including random edited features, image shearing features, two stages edited features and merging edited features with original features in our experiments (more details in Table 1). $L2$ regularization and $L1$ hinge loss are used as the loss function.

Inspired by the bounding-box regression employed in DPM [9] and R-CNN [15], we also train a linear regression model to predict a new detection window using the $5th$ layer features after pooling operation for selective search region proposals, which is to reduce the localization errors.

5.2. Testing Stage

At testing stage, the sliding-window approach and objectness approach are both considered. Although the developments of GPU and parallel technology are rapidly moving forward, computational cost of extracting features from large scale CNNs is still expensive. Precise localization within a sliding-window paradigm is not suitable for combining with large scale CNNs. Thanks to the development of objectness, which has been successful for object detection [2, 32], we apply selective search [32] with "fast" model to generate about 2000 category-independent region proposals for an input image at the test phase like R-CNN [15] and the object detection task is transformed to a standard image classification task. Then non-maximum suppression with 30% threshold is used on the scored windows.

VOC 2010	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
FE_M	**74.8**	**69.2**	**55.7**	**41.9**	**36.1**	**64.7**	**62.3**	69.5	**31.3**	**53.3**	**43.7**	69.9	**64.0**	**71.8**	**60.5**	**32.7**	**63.0**	**44.1**	**63.6**	**56.6**	**56.4**
R-CNN	71.8	65.8	53.0	36.8	35.9	59.7	60.0	**69.9**	27.9	50.6	41.4	**70.0**	62.0	69.0	58.1	29.5	59.4	39.3	61.2	52.4	53.7
Regionlets	65.0	48.9	25.9	24.6	24.5	56.1	54.5	51.2	17.0	28.9	30.2	35.8	40.2	55.7	43.5	14.3	43.9	32.6	54.0	45.9	39.7
SegDPM†	61.4	53.4	25.6	25.2	35.5	51.7	50.6	50.8	19.3	33.8	26.8	40.4	48.3	54.4	47.1	14.8	38.7	35.0	52.8	43.1	40.4
DPM v5†	49.2	53.8	13.1	15.3	35.5	53.4	49.7	27.0	17.2	28.8	14.7	17.8	46.4	51.2	47.7	10.8	34.2	20.7	43.8	38.3	33.4
$\triangle_{R-CNN}$	+3.0	+3.4	+2.7	+5.1	+0.2	+5.0	+2.3	-0.4	+3.4	+2.7	+2.3	-0.1	+2.0	+2.8	+2.4	+3.2	+2.6	+4.8	+2.4	+4.2	+2.7

Table 2. Detection average precision(%) on PASCAL VOC 2010 test. Row 1 shows our experimental results. Row 2-5 show other competition approaches. SegDPM† [12] and DPM v5† [9] use context rescoring. Row 6 shows the differences between FE_M and R-CNN.

VOC 2012	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
FE_M	**74.6**	**69.1**	**54.4**	**39.1**	**33.1**	**65.2**	**62.7**	69.7	**30.8**	**56.0**	**44.6**	**70.0**	**64.4**	**71.1**	**60.2**	**33.3**	**61.3**	**46.4**	**61.7**	**57.8**	**56.3**
R-CNN	71.8	65.8	52.0	34.1	32.6	59.6	60.0	**69.8**	27.6	52.0	41.7	69.6	61.3	68.3	57.8	29.6	57.8	40.9	59.3	54.1	53.3
$\triangle_{R-CNN}$	+2.8	+3.3	+2.4	+5.0	+0.5	+5.6	+2.7	-0.1	+3.2	+4.0	+2.9	+0.4	+3.1	+2.8	+2.4	+3.7	+3.5	+5.5	+2.4	+3.7	+3.0

Table 3. Detection average precision(%) on PASCAL VOC 2012 test. Row 1 shows our experimental results. Row 2 shows R-CNN algorithm results. Row 3 shows the differences between them.

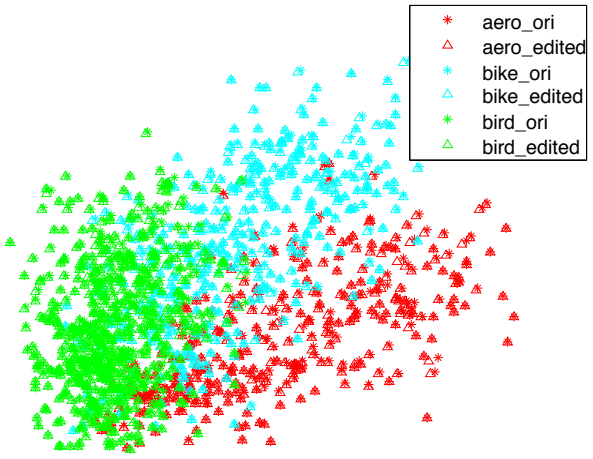

Figure 4. Illustration showing the $5th$ layer CNN feature visualization of our system. The original features and edited features are visualized.

5.3. Feature Visualization

Figure 4 shows the original and edited CNN features extracted from three categories of VOC 2007 trainval set using the $5th$ convolutional layer after pooling operation. In order to display the distribution of high-dimensional CNN features, we apply the principal component analysis (PCA) algorithm [34] to reduce dimensionality and retain two components for visualization, which is inspired by [19]. From Figure 4 we can see that our edited features maintain the distribution of the original features and let the optimal hyperplane more refined which makes the boosted classifier more easy to classify different categories.

5.4. Exp. I: Comparison with Random Edit Algorithm

We train our linear-SVM using train and validation set of VOC 2007. We compare our algorithm with random edit algorithm and the complete evaluation on VOC 2007 test is given in Table 1.

5.4.1 Random Edit Algorithm

$R_i = F_i \otimes W$, $i = 1, 2, \cdots N_C$, where $\otimes$ denotes dot product, W is a random binary vector with 0 or 1 and the length is k. $\frac{m(0)}{m(1)} = \text{threshold}$, and $m(\cdot)$ denotes the number of $(\cdot)$ in the random vector W. And $R_{Edit} = \{R_1, R_2, \cdots, R_{N_C}\}$.

We compare the results of random feature edit with our edit algorithm. Results are shown in Table 1. We find that randomly editing CNN feature units can also obtain competitive performance. This phenomenon proves that our CNN feature edit algorithm is efficient and powerful. Although performance of random feature edit is slightly lower compared to our two stages editing algorithm. The reasons are clear and can be explained between the training set and testing set. The deviation of properties between these two sets do not exist with the same ratio, so editing with propensities is better than random operation.

5.5. Exp. II: Comparison with Image Shearing

We evaluate the approach that enriches training image set by merging sheared images (man-made object parts) and original ones whether can obtain the similar improvement if compared to our editing algorithm. We randomly shear one sub-image from an original object image in the training dataset, and then merge these sheared images into training set. The complete results are given in Table 1. From the re-

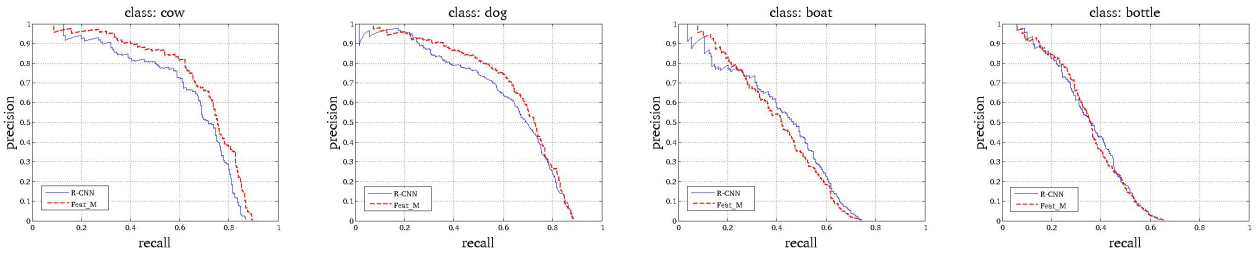

(a) cow (b) dog (c) boat (d) bottle

Figure 5. Four representative categories of Precision/Recall curves on VOC 2007 for R-CNN results compared with Feat_M results. (a) and (b): The categories which our performance is better than R-CNN; (c) and (d): The categories which our performance has a little slight decline compared with R-CNN.

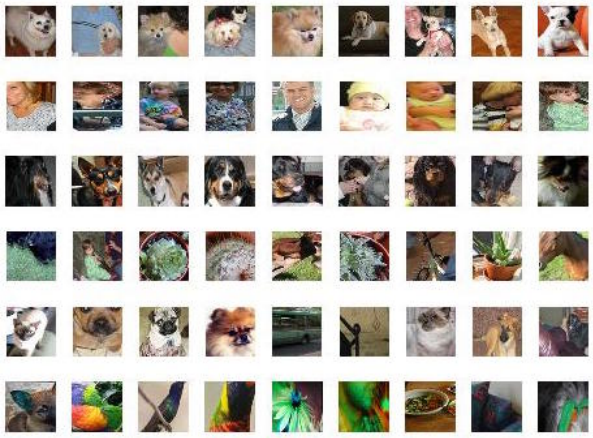

(a) Drops in $Boat$

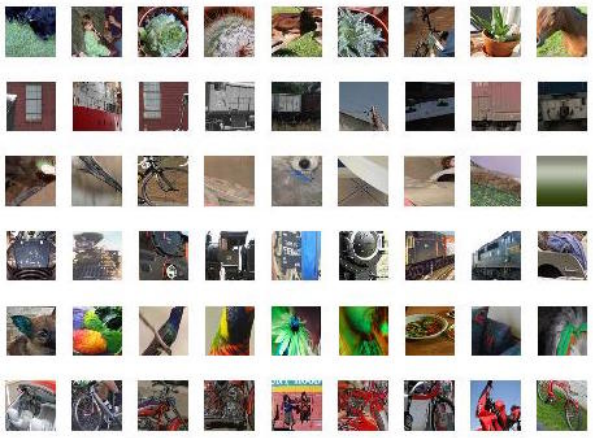

(b) Drops in $Person$

Figure 6. Illustration showing the abandoned channels of $Boat$ and $Person$. These dropped features are helpless for describing their own categories. It's interesting that row 4 and 6 in class $Boat$ also appear in class $Person$ of row 1 and 5.

sults, we can see that merging sheared images is harmful to the whole system. We consider the reasons are as follows: (1) The sub-images are randomly sheared without priori information, if some of these images have no discrimination between them, such as the background, people's chest or the small parts of the sofa. These sub-images will harm the performance of the boosted classifier. (2) The CNN feature units do not simply represent the parts of objects, but also characterize the concepts and properties such as colors, shapes and materials. Editing CNN features is more meaningful than image shearing. Shearing operation could not handle the separation of properties.

5.6. Exp. III: Merge or Not

Simply using edited features and merging with original features for training are both evaluated in our experiments. The performance of merging strategy is improved to 60.1%.

5.6.1 Results Compared with R-CNN

Compared with R-CNN, the average precision of most categories in our algorithm are obviously higher on VOC 2007, especially cow, dog and bus et al. Although a few categories are lower such as boat, bottle, and bike et al. We draw precision recall curves with four representative categories, respectively are cow, dog, boat and bottle showing in Figure 5. For classes cow and dog, our algorithm significantly improves detection performance, but the improvements do not appear in classes boat and bottle. We find that when the object size is small, both the R-CNN and our performance are lower than those with big size. Small objects like bird and bottle are more difficult to be detected, which is also described by Russakovsky et al. [28].

Figure 7 shows the R-CNN detection examples compared to ours. We show the same number of detected windows with top scores in each row. We can see that our method has few false positives (4 $vs.$ 1) in top row, and our false positives have a lower ranking (5, 6) compared to R-CNN (3, 6) in bottom row. These examples tell us that our algorithm can train a better SVM classifier with merged training feature set.

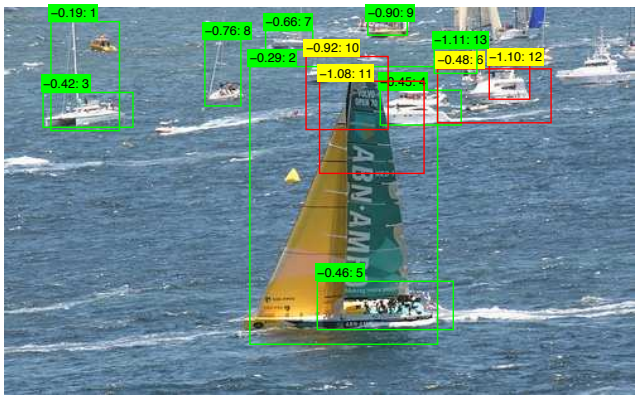

(a) R-CNN

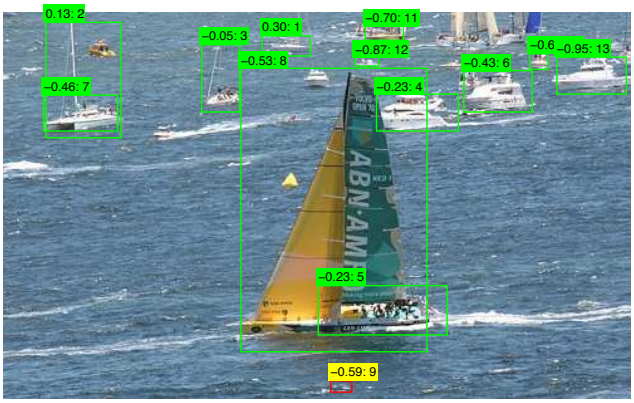

(b) FE_M

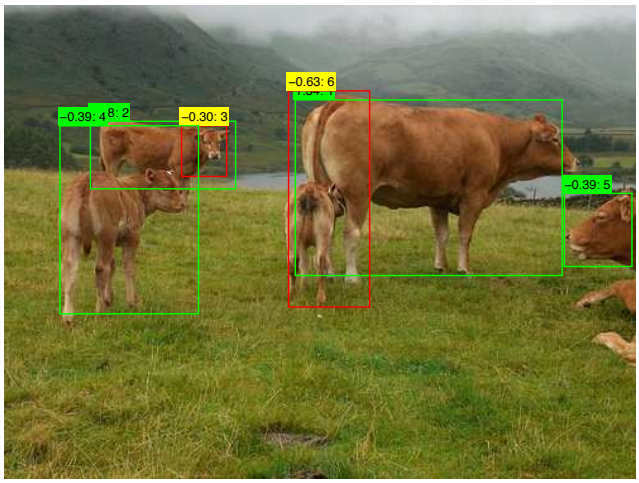

(c) R-CNN

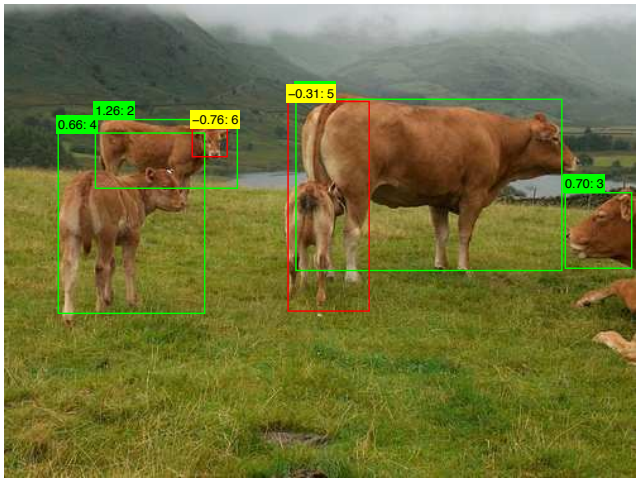

(d) FE_M

Figure 7. Examples on VOC 2007 test set. (a) (c) in the left column are R-CNN detection results, and (b) (d) in the right column are ours. The number $score$ obtained by SVM classifier and the $rank$ of object scores in test image are shown in the top left corner of the object's bounding boxes. Green bounding boxes mean the same size and location of objects detected by both R-CNN and our method, while red bounding boxes mean detection differences appearing between R-CNN and our method.

5.7. Drops Visualization

Figure 6 shows the drops in our algorithm. We visualize these abandoned feature channels by the introduced method in section 2. Two categories are shown in this paper. From this figure we can see that the correspondences of abandoned channels are aligned to those uncorrelated, tanglesome and inutile concepts which are compared to their own categories.

5.8. Results on VOC 2010-2012

The complete evaluation on VOC 2010-2012 test sets is given in Table 2 3. The data sets are identical for VOC 2011 and 2012, so we just present the results of VOC 2012. The results are considerably improved on both of the two datasets. On VOC 2010, we achieve 56.4% mAP vs. 53.7% mAP of R-CNN and on VOC 2012, our performance is 56.3% mAP vs. R-CNN's 53.3% mAP.

6. Conclusions

In this paper we consider the correspondence on a fine feature level than object category and propose an entropy-based model to drop out some negative feature elements in pool_5 to generate somewhat modified CNN features with very low computational complexity. A linear SVM classifier is well trained with merged CNN feature sets. It makes the object detection system to achieve 60.1%, 56.4% and 56.3% performances on PASCAL VOC 2007, 2010 and 2012 test datasets, which are beyond all previously published results. The results indicate that our approach is much more effective with low computational cost.

References

[1] A. L. Berger, V. J. D. Pietra, and S. A. D. Pietra. A maximum entropy approach to natural language processing. *Computational linguistics*, 22(1):39–71, 1996. 3

[2] M.-M. Cheng, Z. Zhang, W.-Y. Lin, and P. H. S. Torr. BING: Binarized normed gradients for objectness estimation at 300fps. In *CVPR*, 2014. 5

[3] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005. 1

[4] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 1

[5] A. Eigenstetter and B. Ommer. Visual recognition using embedded feature selection for curvature self-similarity. In *NIPS*, 2012. 2

[6] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88(2):303–338, 2010. 1

[7] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results. http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html. 4

[8] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2010 (VOC2010) Results. http://www.pascal-network.org/challenges/VOC/voc2010/workshop/index.html. 4

[9] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part based models. *PAMI*, 32(9):1627–1645, 2010. 1, 5, 6

[10] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient matching of pictorial structures. In *CVPR*, 2000. 1

[11] P. F. Felzenszwalb and D. P. Huttenlocher. Pictorial structures for object recognition. *IJCV*, 61(1):55–79, 2005. 1

[12] S. Fidler, R. Mottaghi, A. Yuille, and R. Urtasun. Bottom-up segmentation for top-down detection. In *CVPR*, 2013. 6

[13] E. R. Fischler, Martin A. The representation and matching of pictorial structures. *IEEE Transactions on Computers*, 1973. 1

[14] K. Fukushima. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 1980. 1

[15] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 1, 2, 3, 5

[16] S. Gupta, R. Girshick, P. Arbeláez, and J. Malik. Learning rich features from RGB-D images for object detection and segmentation. In *ECCV*, 2014. 1

[17] B. Hariharan, P. Arbeláez, R. Girshick, and J. Malik. Simultaneous detection and segmentation. In *ECCV*, 2014. 1

[18] K. He, X. Zhang, S. Ren, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. In *ECCV*. 2014. 5

[19] G. Hinton and R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504 – 507, 2006. 6

[20] K. Y. Hui. Direct modeling of complex invariances for visual object features. In *ICML*, 2013. 2

[21] Y. Jia. Caffe: An open source convolutional architecture for fast feature embedding. http://caffe.berkeleyvision.org/, 2013. 5

[22] S. S. H. S. A. K. F.-F. L. Jia Deng, Alex Berg. Imagenet large scale visual recognition competition 2012 (ilsvrc2012). 2012. 1

[23] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 1, 5

[24] Y. Lecun. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1989. 1

[25] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 1

[26] J. Long, N. Zhang, and T. Darrell. Do convnets learn correspondence? *arXiv preprint arXiv:1411.1091*, 2014. 2

[27] W. Ouyang and X. Wang. Joint deep learning for pedestrian detection. In *ICCV*, 2013. 1

[28] O. Russakovsky, J. Deng, Z. Huang, A. C. Berg, and L. Fei-Fei. Detecting avocados to zucchinis: what have we done, and where are we going? In *International Conference on Computer Vision (ICCV)*, 2013. 7

[29] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. In *ICLR*, 2014. 1

[30] K. Sohn and H. Lee. Learning invariant representations with local transformations. In *ICML*, 2012. 2

[31] C. Szegedy, A. Toshev, and D. Erhan. Deep neural networks for object detection. In *NIPS*, pages 2553–2561, 2013. 1, 5

[32] J. Uijlings, K. van de Sande, T. Gevers, and A. Smeulders. Selective search for object recognition. *IJCV*, 2013. 2, 5

[33] X. Wang, M. Yang, S. Zhu, and Y. Lin. Regionlets for generic object detection. In *ICCV*, 2013. 5

[34] S. Wold, K. Esbensen, and P. Geladi. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1):37–52, 1987. 6

[35] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *ECCV*. 2014. 3

[36] N. Zhang, J. Donahue, R. Girshick, and T. Darrell. Part-based R-CNNs for fine-grained category detection. In *ECCV*, 2014. 5

[37] W. Y. Zou, X. Wang, M. Sun, and Y. Lin. Generic object detection with dense neural patterns and regionlets. In *British Machine Vision Conference (BMVC 2014)*, 2014. 5